

Noise and zero point drift in 1.7 μ m cutoff detectors for SNAP

Roger Smith^{*a}, Christopher Bebek^c, Marco Bonati^a, Matthew G. Brown^d, David Cole^b,
Gustavo Rahmer^a, Michael Schubnell^d, Suresh Seshadri^b, Gregory Tarle^d

^aCalifornia Institute of Technology, Pasadena, CA

^bJet Propulsion Laboratory, Pasadena, CA

^cLawrence Berkeley National Laboratory, Berkeley, CA

^dUniversity of Michigan, Ann Arbor, MI

ABSTRACT

We present the results of a detailed study of the noise performance of candidate NIR detectors for the proposed Super-Nova Acceleration Probe. Effects of Fowler sampling depth and frequency, temperature, exposure time, detector material, detector reverse-bias and multiplexer type are quantified. We discuss several tools for determining which sources of low frequency noise are primarily responsible for the sub-optimal noise improvement when multiple sampling. The effectiveness of reference pixel subtraction to mitigate zero point drifts is demonstrated, and the circumstances under which reference pixel subtraction should or should not be applied are examined. Spatial and temporal noise measurements are compared, and a simple method for quantifying the effect of hot pixels and RTS noise on spatial noise is described.

Keywords: HgCdTe, temporal noise, spatial noise, Random Telegraph Signal, noise power spectrum, zero point drift.

1. INTRODUCTION

1.1 SNAP satellite overview

The Super Nova Acceleration Probe (SNAP) is a wide-field telescope, currently in development, which will constrain proposed cosmological models for Dark Energy, using velocity (red-shift) and distance measurements derived from Type Ia Supernovae. A complimentary Weak Lensing survey is also planned which will probe dark matter distribution. Used in combination, these techniques place strong constraints on cosmological models since each method reduces uncertainties in the other.

Control of errors is essential to the success of SNAP. Statistical errors will be reduced by surveying several large areas of sky, while considerable effort is being invested in reducing systematic errors through careful design of the surveys, characterization of the various subsystems and end-to-end simulations. The telescope is configured for maximum stability, and will operate in the thermally stable environment of an L2 orbit.

The two meter class three mirror anastigmat will deliver diffraction limited images to a ~600Mpixel focal plane¹ split equally by area between optical and NIR detectors with *fixed* filters placed close to each. As the telescope scans the sky, objects will be observed in each of 6 optical and 3 NIR wavelength bands. Selected targets can also be directed to a low-dispersion integral-field optical-NIR spectrograph. The optical and NIR detectors will be connected to a common thermal bus, which will be radiatively cooled to at least 140K (130K is possible). Thermal stability will be achieved by operating the detectors at a constant readout cadence for reproducible power dissipation, while the radiators will never see the Earth (or Sun) in normal operation, so

* rsmith@astro.caltech.edu; phone 1-626-395-8780; SPIE paper 6276-30

that only a small amount of power need be dissipated in active control. A separate radiator is planned for the spectroscopic detectors so that they may be operated at a lower temperature to minimize dark current.

Fortuitously, the maximum red shift to be observed only requires supernova observations to $1.7\mu\text{m}$. This cutoff wavelength, results in a sufficiently high bandgap that low dark current can be achieved with radiative cooling. However the noise contribution from the detector material (which exceeds the noise originating in the multiplexor) is greater in HgCdTe tuned for this shorter cutoff wavelength. The SNAP project has funded several rounds of development of $1.7\mu\text{m}$ HgCdTe arrays at both Raytheon and Rockwell. Raytheon, using Liquid Phase Epitaxy to grow the detector material, and Rockwell, using Molecular Beam Epitaxy, have both produced devices, which meet the SNAP noise requirement for imaging. Development continues in order to improve the noise yield and thermal safety margin for the imager, and to meet the more demanding needs of the spectrograph.

1.2 SNAP noise requirements


Exposures for the imaging survey will be 300s long, with the shutter closing for 30 seconds to accommodate readout of the optical detectors. The total noise in a dark 300s exposure is not to exceed 8 electrons. Multiple sampling will be employed while the shutter is closed. Although extending the sampling into the exposure is permissible, this will result in a reduction in effective exposure time. Simulations of the observations and full data reduction have been performed to determine the relative impact of changes in detector QE (or effective exposure time) and noise on errors in the final science products.² These studies show that for the supernova light curves, QE has more leverage than noise on several science parameters. This, combined with the weak improvement in noise observed for deeper sampling, it seems unlikely that sampling during the exposure will be beneficial, so the $8e^-$ noise specification will most likely apply to the case where Fowler sampling is performed for 15 seconds before and after the shutter opens for 300s.

At $8e^-$ the imager noise becomes insignificant compared to the photon shot noise from the zodiacal light in 300s. The spectrometer has one tenth of the imager's spectral bandpass per pixel, and taking optical efficiency into account, sees $87e^-/\text{hr}/\text{pixel}$ from the zodiacal light, so that shot noise will be degraded 50% by a total dark noise of $7e^-$ in a half hour exposure. We present data showing that this is currently achievable.

1.3 Test system

Detector controllers from Astronomical Research Cameras Inc. (often referred to as "Leach III controllers") were used for both Rockwell HAWAII-2RG (H2RG) and Raytheon VIRGO tests. The bias voltage generators on the 8 channel video boards were used, modified by the addition of substantially greater filter capacitances and much lower (20ohm) series output resistors than shipped. In tests with bias voltages looped back to the video input, or with input grounded, it was found that bias and offset generators drifted significantly ($\sim 1\mu\text{V rms} = 1e^-$) in 300s. An alternate bias generator board developed by the Australian National University was found to have excellent drift characteristics providing a point of comparison for drift suppression algorithms utilizing reference pixels. (see section 3.4)

Fourier transforms of time series data in the bias loopback configuration revealed strong harmonics of the ARC Inc. switching power supply, which were not entirely suppressed by the additional filtering. It proved cheaper to replace the supply with an assemblage of linear power supply modules than to track down the cause. The residual harmonic power was negligible (see Figure 8).

Analog, digital, and shield ground s were joined on the backplane of the controller. Each video signal was paired with analog ground from the detector mount, which was fed to the reference input of the differential preamp. In future we plan to bring Dsub from the detector mount to the reference input in order to suppress Dsub drift, which is superimposed on the video signal, particularly since coupling within the H2RG mux of ground noise to the signal is weak.

A Lakeshore Model 340 temperature controller provided 1mK readout resolution and comparable servo control accuracy for the Molybdenum plate on which the detector was mounted. A calibrated RTD was used in place of a forward biased diode for temperature sensing so that transients induced by clock pickup would not be rectified and thus have no impact on the slow thermal control loop. Thermal slew rate was held precisely to 0.5K/min at all times through manipulation of a solenoid valve on the liquid nitrogen tank, detector thermal servo, and a 3rd thermal servo on the N₂ tank worksurface.

Detectors were mounted in a windowless light tight chamber sealed with labyrinths. Flex circuit and enameled wire feedthroughs were potted in silver filled epoxy to provide a light seal and thermal contact. Illumination was provided by switched NIR LEDs inside this dark chamber. For flux stabilization, LEDs were thermally regulated to 1 mK rms, and powered by precision current sources. A metal mask placed <1mm from the detector surface covered ~20% of the detector area so that any light leak would appear as a shadow on the detector. None were detected in 10 hour dark exposures. After dark current tests were performed, a neutral density filter was attached to the mask to allow the LEDs to operate at 1000 times the flux (~1mA drive current) where their flux is more stable.

2. NOISE MEASUREMENT METHODS

2.1 Spatial noise

Pixel-to-pixel variance in a single dark frame is frequently used to estimate the read noise, which is strictly the frame-to-frame variance for a given pixel. The assumption here is that the noise varies little from pixel to pixel so that the average read noise is an adequate descriptor. An accurate estimate of the standard deviation is easily achieved by simply including many pixels in the statistics box (eg. 100x100). We confirmed the low statistical uncertainty by comparing many such statistic boxes across the image area.

Spatial noise is also popular since it is insensitive to common mode noise, which can be produced by bias voltage and temperature changes and tend to become increasingly problematic in long exposures. Very low frequency noise originating in the output amplifier and video chain will appear as spatially correlated shifts, contributing to the spatial noise only if the statistics box extends across multiple amplifiers. Thus the spatial noise should be considered a measure of the average noise *achievable* when environmental factors are well controlled or corrected

Spatial noise can be degraded by fixed patterns, which may be caused by differences from pixel to pixel in offset drift due to self heating/cooling during the exposure, or by spatial variations in dark current. Reproducible effects are removed by taking pairs of identical dark frames, computing the standard deviation of the difference frame, then dividing by $\sqrt{2}$. However this does not remove the shot noise from hot pixels. The effectiveness of frame differencing in suppressing dark current structure, is demonstrated in Figure 1. Pixels with dark current above some threshold were excluded in the standard deviation of single frames. As the dark current threshold is lowered, the noise approaches that of the difference frames. The fact that excluding all but the lowest dark current pixels does not produce a *lower* noise estimate than the frame differencing, indicating that the shot noise from the high dark current pixels is not biasing the estimate of average noise significantly for difference frames.




Figure 1: Comparison of single-frame spatial noise when pixels are excluded on the basis of dark current, with spatial noise of the difference frame. As the maximum acceptable dark current is lowered the results converge on the frame differencing case but never fall below it, demonstrating that while the fixed dark current pattern degrades spatial noise, the shot noise of the dark current is (in this case) in significant. As expected, the effect is more pronounced at 140K (left) than at 130K (right).

2.2 Temporal noise

It is “temporal noise” which the user experiences: i.e. fluctuations in a given pixel value when an exposure is repeated. This includes electronic and thermal drifts without regard for their origin. Our primary goal was to monitor and sometimes steer the development of lower noise detectors, for which spatial noise is generally adequate. However, only temporal noise measurements yield spatial noise maps and histograms. We also study temporal noise issues in order to develop electronic drift and thermal control requirements for flight electronics, and to evaluate reference pixels subtraction algorithms(etc) which may relax these requirements.

The standard deviation of a finite sample set contains both random and systematic errors when used to estimate the standard deviation of the parent population. The systematic error is corrected by multiplying by $\sqrt{n/(n-1)}$, where n is the number of samples. The random error scales as $1/\sqrt{n}$. While thousands of exposures are required to obtain precision estimates of temporal noise independently for each pixel the *mean* noise can be estimated accurately by combining many less precise temporal noise estimates. Taken to the extreme, the mean temporal standard deviation from just two frames, when multiplied by $\sqrt{n/(n-1)} = \sqrt{2}$ equals the standard deviation of the difference frame divided by $\sqrt{2}$. In practice we have used 20 to 200 frames to estimate the temporal noise, avoiding lower frame numbers so that the common mode drifts will be adequately sampled, and avoiding high frame counts when exposure times are long.




Figure 2: The effect of number of exposures on temporal noise histogram shape (1.3s exposures, Fowler-1).

With few samples, histograms of these temporal noise maps broaden and assume a normal distribution with a peak, which under estimates the true mean noise. As the number of images to make the noise maps is increased the true histogram shape, width and modal noise is revealed. In a log linear plot the Normal Distribution is an inverted parabola, and any excess power in the wings is readily apparent. The left most peak is produced by the reference pixels. Figure 9 shows histograms for a different detector and much longer exposure time.

Our goal was to demonstrate close agreement between temporal noise and spatial noise for a broad range of exposure times and sampling depths, since this provides a convincing demonstration that the noise measured is both real and detector limited. To achieve this agreement in long exposures, as shown in Figure 6, considerable effort was invested in controlling zero point drifts, as discussed below.

2.3 Gain

The inverse system gain (e-/ADU) was derived from pairs of matched flat fields with intensity from zero to only $\sim 10,000$ e- to avoid non-linearities. The spatial variance of difference frames was plotted against the sum of the means, then inverse slope was multiplied by the correction factor, 0.87, derived from the shot noise auto-correlation function (As described by Andrew Moore^{4,5}) to compensate for the attenuation of the variance by inter-pixel capacitance.⁶ The corrected inverse gain was 1.13e-/ADU so that typical single sample detector noise was ~ 15 ADU, and thus always much greater than the 1.1 ADU noise with video inputs shorted ($3\mu\text{s}$ integrator dwell time). This is not to say that electronic noise sources were insignificant contributors to drifts after sample averaging.

3. ZERO POINT DRIFT

3.1 Sources of zero point drift

Charge traps in near the MOSFET channel induce bimodality into the voltage distribution of the pixel buffer and output source followers, called “Random Telegraph Signal”. The random timing of electron capture and emission results in a $1/f^2$ noise power spectral density at low frequencies while the superposition of several such traps, of random amplitude and time constant produces a $1/f$ noise spectrum, which can extend to very low frequencies.³ Ideally this low frequency noise power at the pixel and output amplifier would be the only sources of zero point drift during an exposure. Unfortunately we find that noise sources with similar spectral characteristics exist in the detector material itself, which dominate the multiplexor transistor noise.

Table 1 shows that several bias voltages couple strongly to the video output of the H2RG multiplexor, and while the VIRGO multiplexor is less sensitive it is not totally immune. Bias noise on timescales less than 100ms can generally be attenuated capacitively. However, temporal noise can be seriously degraded by bias voltage drifts on timescales beyond those which can be filtered. Since Bias voltage noise on sub-frame timescales is attenuated efficiently by multiple sampling, whether Fowler or least squares fit: it is the low frequency noise or bias drift, which is most problematic.

Table 1: Output voltage sensitivities to bias voltage fluctuations measured for a bare mux. Actual values are expected to be slightly lower when the detector capacitance is present. Rockwell’s simulations⁹ predict 10-20% lower numbers. We confirmed Rockwell’s prediction of a flat frequency response to about 500kHz.

H2RG $\delta\text{Output}/\delta\text{Bias}$		VIRGO $\delta\text{Output}/\delta\text{Bias}$	
Dsub	TBD	Vdetcom	TBD
BiasPower	+1.7	VPuc	0.7
BiasGate	-1.5	VNout	0.06
VddA	-0.4	Islew, VNuc	0.05

Others	#0.04	Others	#0.03
--------	-------	--------	-------

3.2 Thermal effects

The gate-source voltage of both the pixel buffer and output amplifier are quite sensitive to temperature (we measure 102e-/K for H2RG-103, Hall⁷ measures ???]. Approximately ???%rms pixel-to-pixel variation in temperature coefficient is measured by Hall⁷. Thus thermal drift will produce both common mode, channel to channel, and pixel-to-pixel signal variations, which can degrade temporal noise and, to a lesser extent, spatial noise. In the lab we have controlled detector temperature to 1.7mK rms, implying that <0.18e- zero point drift derives from thermal stability. This would be an underestimate if temperature sensing errors occurred slowly enough for the servo to introduce a compensating temperature variation: the stability of the reading alone does not prove that the temperature is stable. However, we saw no correlation between the mean value in the image, the temperature reading or the servo power implying that sensor errors do not account for the residual zero point drift.

3.3 Reset Anomaly

The HAWAII-2RG and VIRGO utilize 3 transistor unit cells, which only dissipate power when addressed. Similarly the CMOS logic at the edge of the multiplexor only dissipates power during clock transitions. The pixel power dissipation point rasters across the multiplexor so that the pixel temperature is always varying slightly. The key is to clock the device so that this thermal variation reproduces accurately from scan to scan: if the clocking cadence varies, the consequent temperature change can produce large offset voltage drift during the exposure. This effect, once thought to be a transient induced by reset and dubbed “Reset Anomaly” is now known to be self heating (and cooling). For some clocking schemes such as “global reset” a slope can be introduced within a frame in addition to a common mode drift. If the transition from “idle” to “read” involves aborting a raster scan part way through, then the thermally induced drift can vary and thus contribute to temporal noise.

Reference pixel subtraction fails to correct for self-heating variations due to thermal gradients from center to the edges where the reference pixels are located. Thus the ideal arrangement is to idle, reset and sample up the ramp at a constant cadence, with no changes in frame rate or truncated scans at the start of the exposure. The H2RG exhibits no measurable multiplexor glow, so sampling can continue up the ramp whether or not the central samples are used, in order to avoid cooling during the exposure delay. In practice we find that ~3 dummy read scans (at 1.3s each) prior to the second sample group is sufficient to bring the H2RG to thermal equilibrium due to self-heating. This has been employed for the noise data presented in this paper.

By contrast the 2.5μm VIRGO array we tested⁸ can take several minutes to reach thermal equilibrium after a major clock cadence change and suffers from mux glow of 0.04±0.01e-/read. This glow is expected to depend on time spent addressing each pixel and mux biasing. VIRGO mux glow in 1.7μm material and dependencies on biasing and clocking have not been studied yet, as we await the imminent delivery of devices with lower dark current and hot pixel count.

3.4 Reference pixel subtraction

To determine the optimum reference pixel subtraction scheme, we acquired ~200 dark exposures of 300s duration at various Fowler sampling depths, then measured the correlation between the means taken over various image and reference pixel areas. The fitted slope indicates that the full reference pixel deviation should be applied to the image area, as expected. The rms deviation orthogonal to this slope indicates the uncorrelated noise, which is added in the process of correcting for the correlated component. Whether and how this reference pixel subtraction should be applied depends on the strength of the correlation, which is related to the source of the drift.




Figure 3: Inappropriate reference pixel subtraction: channel 32 reference columns were averaged and subtracted from all other channels for fowler-32 sampling with the ANU bias board which exhibited very low drift, so that the dominant source of drift was the input offset of the 8ch ARC Inc video card. While image pixels on the same channel as the reference pixels clearly benefited, the other channels on the same video card only benefited partially. Channels read through different video cards were degraded by reference pixel subtraction since the drifts were uncorrelated.

For CDS frames, which exhibit significant drift on sub-frame timescales (banding) lateral reference pixel subtraction improves the temporal noise histograms. For the deeper fowler sampling of interest to SNAP this faster drift component rapidly averages away so that the dominated noise sources correlate across multiple frames. In this case the lateral reference pixel subtraction does more harm than good. The long correlation time ensures that the top and bottom reference pixels correlate well with the image pixels even though they are not measured concurrently. Top and bottom reference pixels for given channel should be averaged together for maximal noise reduction, then subtracted from image pixels read through the same channel. This suppresses drift in video input offsets, 1/f noise in the multiplexor output amplifier, thermally induced drifts in the mux output, and common mode bias drift.




Figure 4: Temporal noise histograms with (thick lines) and without (thin lines) per-channel top-bottom reference pixel subtraction. The ANU bias board has sufficiently low drift that no improvement is seen from reference pixel subtraction in CDS (Fowler-1) frames, however at increased sample depth the random effects are averaged away so that correlated drifts (probably video input offset drift rather than bias drift) represent a greater proportion of the noise. The larger bias drifts originating in the ARC Inc (Leach) video board are also substantially corrected by top bottom reference pixel subtraction.

For Fowler-32 sampling in 300s exposures, the scatter in the global mean is $5.8e^-$ for the ARC (Leach) biases, reducing to $1.1e^-$ after per-channel top-bottom reference pixel subtraction. The ANU bias board is much more stable delivering less scatter (0.7ADU rms) in the uncorrected global mean. This reduces to just $0.2e^-$

rms after reference pixel subtraction. Although the reference pixel subtraction attenuates the zero point drift by a factor of 4 or 5, it is clearly beneficial to have stable biases in the first place. Further work is planned to measure the attenuation more directly.




Figure 5: Scatter in mean and spatial standard deviation, (at left) without and (at right) with per-channel top-bottom reference pixel subtraction, for various fowler sampling depths at $1.1e-/ADU$. The scatter in the mean zero point for the whole frame is reduced substantially. The first frame of the sequence was subtracted from each of the others to remove fixed pattern structure before calculating the standard deviation.

4. AGREEMENT BETWEEN TEMPORAL AND SPATIAL ESTIMATES

Figure 6 shows that excellent agreement has been obtained between spatial and temporal noise, even at very long exposure times. Since many days were required to acquire the temporal noise data, Figure 6 also implies the noise is very stable on that timescale. To achieve this agreement, all of the following were required:

- Good thermal stability at the detector mounting plate, due to the strong dependence of offset on temperature.
- Several dummy reads before each fowler sample group, or continuous constant cadence clocking, to eliminate the effect of variable self-heating.
- Frame differencing for the spatial noise, to suppress fixed patterns primarily due to dark current structure.
- Top-bottom reference pixel subtraction, channel-by-channel, to suppress the effects of bias voltage drifts, $1/f$ noise and offset drifts in the output buffer and video chain.
- Scaling of the mean temporal variance by $\sqrt{N/(N-1)}$ to estimate the population variance. No such scaling was needed for the spatial noise due to the large number of pixels included in the spatial noise estimate.




Figure 6: When the appropriate measures were taken, the mean temporal noise agreed with the spatial noise in long exposures to better than 1%, indicating that both measurements are probably valid and that temporal noise maps are not corrupted by drifts.

5. DETECTOR NOISE DOMINATES MUX

Noise versus number of fowler samples has been measured for a bare H2RG multiplexor and shown to be similar to the noise measured when the reset switch is held permanently closed throughout the readout. In this configuration the Vreset noise is added in quadrature with the noises contributed by the multiplexor itself, BiasGate, BiasPower and Vdda which couple significantly to the video output. In fact these biases couple so strongly that their noise can easily dominate the noise generated in the multiplexor itself.

Root square subtraction of the noise with permanent reset from the full system noise seen in normal operation places a lower bound on the noise from the detector alone. In Figure 7 though Figure 11 it is clear that the noise originating in detector material totally dominates multiplexor and controller noises.

This noise differential increases with increasing fowler depth and increasing exposure time since the multiplexor noise averages away more efficiently than detector noise, indicating that the detector material exhibits more noise power at low frequencies. An examination of the noise on the bias voltages when looped directly to the video input (bypassing the mux) shows considerable very low frequency noise power (drift) for biases derived from the 8ch video card from ARC Inc. However the use of spatial noise suppresses bias voltage and video drifts, along with low noise frequencies in the mux output amplifier, which, at low frequencies, correlate across the entire frame. In the case of temporal noise these are suppressed by top-bottom reference pixel subtraction.




Figure 7: Comparison of the total H2RG-103 detector noise with bare multiplexor (#006), with individual bias voltage looped back directly to the video input, and with a shorted video input. While the combined effect of the various bias voltage noises comprises a significant fraction of the noise seen from the bare mux, this contributes only a few percent to the total noise. The open squares show that the noise with detector permanently reset is slightly greater than the bare mux noise since the noise on Vreset (comparable to BiasGate noise) is fully coupled to the output.

6. NOISE POWER SPECTRAL DENSITY

The noise power spectral density (PSD) for five widely spaced but otherwise randomly selected pixels is shown in Figure 8, at left. This was obtained by stopping the clocks at a given address, waiting three minutes for the pixel to thermally equilibrate, then reading at the normal data rate (100kHz) and video bandwidth, without clocking the multiplexor for 160 seconds. This was repeated (with 3s gaps) 18 times. The average frame was then subtracted to remove substantial residual thermal drift and the (almost negligible) dark current ramp. Each frame was then Fourier transformed to produce power spectra, which were averaged. The process was performed at five widely spaced pixel addresses. At two of these addresses the process was repeated with reset switch held closed to produce the power spectrum of multiplexor and system noise. Total noise and multiplexor noise spectra are compared in Figure 8, at right. The relative contribution of the multiplexor to total noise is markedly different in each case, but still secondary to detector noise. Interference harmonics are prominent at the high resolution obtained from the long time series, however their contribution to the total power is negligible. Note that linear power supplies were used during these tests.

Noise power is dominated by the lowest frequencies. Given the single pole roll off (rather than a lower power), Fowler sampling would not be expected to be very effective, but we note that the most important frequencies for fowler sampling performance are below those reached here. In future we plan to derive (lower resolution) spectra for many pixels from long sequences of sample up the ramp data using standard clocking so that self-heating will be constant, and to extend this noise power spectrum down to the sub-millihertz range.




Figure 8: At left, Power Spectral Density for widely separated pixels chosen at random from H2RG-103 at 140K and 0.5V reverse bias. All pixels were read through the same channel, and video chain. At right, two of these pixels were re-measured with reset switch closed to show the noise contribution from multiplexor and controller. The higher mux noise corresponds to the pixel with lower total noise. The turn-up below 0.1Hz is due to imperfect subtraction of offset drift due to pixel buffer transistor self heating. The frequencies of interest for predicting Fowler sampling performance in 300s exposures are below this range.

7. NOISE DISTRIBUTIONS

With drifts under control, it becomes possible to measure the pixel-to-pixel noise distribution. A large number of exposures must be acquired to gain sufficient statistical accuracy in the estimate of the variance. This must be done at the exposure time and sampling depth intended for the observational program. Figure 9 shows the histogram for sequences of 300 exposures each of 300s duration (25 hours total exposure time per curve): the systematic underestimate of the mean standard deviation due to the finite number of samples is only $\sqrt{300/(300-1)} - 1 = 0.2\%$. The Full Width Half Maximum (FWHM) of the noise distribution (13%) is dominated by the statistical uncertainty in the estimate of the standard deviation ($2/\sqrt{300} = 11.6\%$). Subtracting the statistically induced width in quadrature from the measured width provides a rough estimate of the true FWHM of the noise distribution, 7%.




Figure 9: Temporal noise histograms for 300 exposures, each of 300s duration, for Fowler 1 (CDS) and Fowler 16 sampling. For the dotted histograms, the reset switch was held closed to short out the detector noise providing a measure of the of multiplexor/controller noise.

Figure 9 illustrates how Fowler sampling not only reduces the mean, but also reduced the width of the distribution. While the core of this noise distribution is very compact, it exhibits a substantial high tail. Nine percent of pixels lie more than three-sigma above the mode (compared to 0.15% for a Normal Distribution). Four percent of pixels exhibit more than twice the modal noise, while three percent are more than three times the mode. This high tail is more pronounced for the permanent reset case. Continuous sampling of the high noise pixels on a bare mux or with permanent reset reveals that the high noise has a bimodal characteristic called Random Telegraph Signal (RTS) which is indicative of a single trap in the channel of the pixel buffer MOSFET³, and while common to all MOSFETs, is exacerbated by small geometries.




Figure 10: Pixels on a bare mux (or a detector with reset switch permanently closed), which exhibit high noise, do so because of (usually) single traps in or near the buffer MOSFET channel. The capture and release of an electron produces a bimodal output voltage called Random Telegraph Signal. This plot shows the output of a single noisy pixel at two different temperatures and markedly different temperature dependence of capture and emission time constants, while amplitude which is primarily a function of trap location remains constant as expected.³ Each units on the x axis is equivalent to 640 μ s for a full span of ~700 milliseconds.

Figure 11 shows that root square subtraction of the multiplexor-only temporal noise map from the total noise map yields a detector noise distribution, which is more Gaussian (at the level of statistical precision achieved with 300 frames for each noise map). Analysis of the corresponding cumulative histograms (not shown) reveals that of the 9% of pixels lying more than three sigma above the modal noise, 1% are due to high dark current, 3% are due to high multiplexor noise (presumably RTS noise) and 4% are due to high detector noise. Of the 4% of pixels with than twice the modal noise, half are due to multiplexor noise, a quarter are due to detector noise and the other quarter (1% of all pixels) are due to are due to high detector noise: mux RTS noise and dark current are the primary causes of pixel inoperability.




Figure 11: Root square subtraction of the multiplexor noise map (obtained by permanently closing the reset switch) from the total (normal) noise map, yields a more Gaussian distribution, an inverted parabola in the log-linear plot at left. Only a moderate fraction of the high noise tail is removed by excluding pixels with greater than mean dark current. The log-log plot at right covers an expanded range and demonstrates that the high noise tails follow a power law rather than exponential distribution.




Figure 12: Noise vs. total sampling time for various sample rates, for H2RG-103, 0.25V reverse bias, at 110K. When sampling at half the usual frame rate, video bandwidth was held constant and a delay was inserted between frames. To achieve twice the frame rate only half the frame was readout (so pixel rate was unchanged).

Figure 12 shows that at high enough sampling depth the number of samples is no longer as important as the time spent sampling. This is to be expected since the white noise component (uncorrelated from frame to frame) averages away until the low frequency noise dominates. In this condition, noise is correlated across several frame times so there is no advantage to sampling more frequently, and little penalty sampling more slowly. This will ease the data rate requirement for the SNAP spectrograph where longer exposures and deeper Fowler sampling are to be employed.

8. TEMPERATURE DEPENDENCE




Figure 13: H2RG-038 detector noise for 300s CDS exposures (left) is highly dependent on temperature. This is not accounted for by dark current, which is similar to H2RG-103 (Figure 14). By contrast H2RG-103 (right), which was made using a different recipe (currently the SNAP baseline), shows essentially no dependence of read noise on temperature. The slight temperature dependence of total noise seen here can be explained by dark current, which begins to have some effect in the longer (1000s) exposure time for the plots at right.

9. DARK SIGNALS

H2RG-103 is typical of the Rockwell's 1.7 μ m cutoff HgCdTe, which reliably achieves sufficiently low enough dark current at the 140K upper temperature limit so that impact on mean noise in 300s exposures is minor. Operation at 130K is believed possible, at which point shot noise from dark current is negligible. As demonstrated in Figure 1, the subtraction of dark current is still necessary to removed fixed patterns. Integration times up to an hour are being contemplated for the spectrograph, which is cooled with a separate radiator. Figure 14 shows that the dark current floor is reached at 120K for H2RG-103, and Figure 6 shows that once a matched dark frame has been subtracted, the spatial noise is unaffected by dark current even in 300s exposures.






Figure 14: Modal dark current scales according to the theoretical curve down to 135K then approaches a $0.014e^-/s$ floor at 0.25V reverse bias. The scaling implies a bandgap, which is very close to hc/λ_{cutoff} .

The noise and dark current performance are seriously compromised both by prompt and delayed optical emissions from the CdZnTe substrate caused by interactions with a variety of particles. Figure 15 compares the Cosmic Ray events in long exposures for substrate-removed and substrate-intact devices. The substrate is ~ 30 times thicker than the pixel is wide, so particle interactions high in the substrate deliver light over a wide area. When the substrate is removed particle events usually only affect single pixels since the HgCdTe layer is thinner than the pixel width.




Figure 15: Cosmic ray events with and without CdZnTe substrate removed from H2RG detectors. Each image is the difference of consecutive dark frames so dark current and hot pixels are suppressed. Particle events can be positive or negative. *Left*: substrate intact, 300s exposures, $1.7\mu m$ H2RG-038 @100K; *Center*: Substrate intact, 3000s exposures, $2.5\mu m$ H2RG-019 @80K; *Right*: substrate removed, 3000s exposures, $1.7\mu m$ H2RG-019 @110K

10. EXPOSURE TIME DEPENDENCE

Figure 16 shows the noise reduction with Fowler sampling depth for constant effective exposure time, which is defined as the exposure between the first sample of the first group of Fowler samples and the first sample of the second group, and represents the mean signal integration time which in turn determines the photon shot noise. If the noise spectrum was white, the read noise (not shot noise) would reduce as $1/\sqrt{N}$. However the

detector material, like most semiconductor components exhibits increasing noise at low frequencies, so the noise improvement scales as a lower power of the number of samples.




Figure 16: Left: Noise versus Fowler Sampling depth for various exposure times. (Exposure delay is adjusted for Fowler sampling depth to maintain constant effective exposure time). Mean dark current (see Figure 14) is insufficient to have much impact on mean noise. Flattening of these curves must primarily be due to the low frequency noise properties since exclusion of high dark current pixels (not shown) has little effect on noise.

For a given Fowler sampling depth, noise rises with exposure time, due to the very low noise components often referred to as “drift”. Very low frequency noise is correlated from frame to frame causing the fowler noise curves to flatten, however more time is available for Fowler sampling. For H2RG-103, which represents the current baseline recipe for SNAP, the noise increase with exposure time is fully compensated by deeper fowler sampling. Figure 1 shows noise verses exposure efficiency, which we define as “effective exposure time (a function of Fowler sampling depth) divided by total time per frame”. The curves for different exposure times nearly overlay: while noise is strongly dependent on exposure efficiency, it is only weekly dependent on total time per exposure.




Figure 17: Noise vs. exposure duty cycle (a function of sampling depth) varies slowly with total time per exposure but is responsive to higher reverse bias, particularly for the long exposure times. Since these exposure times are under consideration for the SNAP spectrograph data are shown here for lower temperature operation.

Figure 18: Thick lines show total noise (left axis) for various fluxes and hypothetical read noise dependence on number of fowler samples (and thus exposure duty cycle). Fine lines show the corresponding signal to noise ratio (right axis). Maximum S/N occurs at sampling depth where the tangent to the total noise curve passes through the origin

11. OPTIMAL FOWLER SAMPLING DEPTH

As the number of Fowler samples, n , is increased the exposure efficiency, E , reduces as,

$$E \equiv \frac{T - nt_s}{T} \quad (1)$$

where T is the total time per exposure and t_s is the time per sample. The noise reduction obtained has to be balanced against the resulting loss of signal. For example, if, for constant total exposure time, T , noise scales as

$$N(n) = N_1 n^{-\alpha} \quad (2)$$

then, for the read noise dominated case, signal-to-noise ratio, R , is maximized when

$$0 = \frac{dR}{dn} = \frac{d}{dn} \left[\frac{(T - nt_s)f}{N_1 n^{-\alpha}} \right] \quad (f = \text{flux}) \quad (3)$$

which reduces to

$$nt_s = \frac{\alpha}{1 + \alpha} T \quad \text{and} \quad E \equiv \frac{T - nt_s}{T} = \frac{1}{1 + \alpha} \quad (4)$$

For white noise $\alpha=1/2$, and we have the familiar result that each group of Fowler samples should occupy one third of the total time and $E=2/3$. Equation (4) tells us that, as the noise becomes a weaker function of n due to the presence of stronger low frequency noise, the optimum signal to noise occurs with fewer fowler samples and higher exposure duty cycle. This is only quantifying what is intuitively obvious: we are less willing to give up exposure efficiency (signal) when deeper sampling yields less noise improvement.

A simple graphical technique can be used to find optimum sampling depth for a given exposure time. Total noise from all sources is measured for the constant total exposure time of interest (ie. T not the effective exposure time $T - nt_s$). Armed with the empirical total dark noise expressed as a function of effective exposure time, $N_T(E)$, we can write the signal to noise definition:

$$R = \frac{fET}{N_T(E)} \quad (5)$$

Differentiating to find the exposure efficiency for maximum signal-to-noise

$$0 = \frac{dR}{dE} = \frac{fTN_T(E) - fET \frac{dN_T}{dE}}{N_T^2(E)}$$

$$\frac{dN_T(E)}{dE} = \frac{N_T(E)}{E} \quad (6)$$

)

Thus, highest signal to noise occurs where the tangent to $N_T(E)$ passes through the origin. The optimum exposure efficiency (and thus sampling depth) can be obtained for any signal level, S, by plotting

$\sqrt{S \cdot E + N_T^2(E)}$ then graphically finding the tangent which passes through the origin. Figure 18 shows total noise and signal-to-noise curves for the hypothetical but not atypical case where noise scales as $n^{0.25}$.

Returning to Figure 17 we see that the optimum Fowler sampling depth for read noise limited pixels would be at an exposure efficiency, $E=80\%$ for 0.25V detector reverse bias and $E=85\%$ for 1V bias. Pixels with greater signal must be sampled fewer times if S/N is to be optimized. Current plans for the spectroscopic detector (only) allow for storing and transmitting all data points sampled up the ramp, so each pixel may be optimally post processing to obtain the best S/N depending on flux detected. Least squares fit or other optimal filtering algorithms have yet to be tested. We note however that for fowler sampling the optimum depth is only a soft function of intensity so it may be possible to adopt a fixed algorithm with little sacrifice.

12. DETECTOR BIAS

Figure 17 shows noise for typical detector reverse bias (0.25V) and an extreme value (1.0V), for various exposure times. The conversion gain was measured at both biases, decreasing from 1.23e-/ADU at 0.25V bias to 1.08e-/ADU at 1V bias. The correction for inter pixel capacitance remained constant. Some of the total input a capacitance is due to the multiplexor and bump bonds, which tends lower the percentage change due to the detector bias. In all cases the noise is clearly lower at higher reverse bias.

The change in noise correlates with the change in capacitance in the case of short exposures, suggesting that the detector noise is a voltage source rather than a charge source, but there are at least two reasons why this does not appear to be the correct explanation. While the noise source would be expected to be equivalent to an input *voltage* source, if the multiplexor was the dominant source of noise, mux noise has been shown to be a minor component. If the noise was correctly modeled as an invariant voltage source, converted to an equivalent charge noise by a detector bias dependent capacitance, then the scaling with detector bias should be *invariant* with exposure time, but this is not so: we find that the reduction in noise with increasing reverse bias is significantly greater for long exposure times. (The down side is a twofold increase in dark current.)

It seems more likely that the detector noise is due to charge trapping and release. Trapping does not induce a voltage change per se, but could modify the charge density profile and thus the electric field across the diode junction. There is no charge available to be captured in the depletion region, so it is plausible that the explanation for lower noise at higher bias is that traps are disabled in the newly depleted material. The greater impact on long exposures would imply that the longer time constant traps are concentrated closer to the diode junction than short time constant traps and thus are preferentially disabled as depletion width increases.

13. DETECTOR COMPARISONS

QE, intrapixel response and interpixel capacitance have been reported in a companion paper¹ for HgCdTe grown on CdZnTe substrates by Liquid Phase Epitaxy at Raytheon and have been compared with material grown by Molecular Beam Epitaxy at Rockwell. This paper has focused primarily on the Rockwell devices since the noise in the Raytheon arrays are compromised by high dark current (mean = 0.18e-/s at 140K) and a high incidence of hot pixels. This topic will receive more attention soon after this conference when the product of the current R&D effort at Raytheon is to be delivered. Meanwhile

Figure 19 provides a preview of what might be expected.




Figure 19: At left, noise versus sample depth at the SNAP imaging exposure time and upper temperature for three candidate technologies. The noise in the Raytheon VIRGO-2K (V2K) and Goodrich InGaAs (on H2RG mux) is degraded by high mean dark current even when (in the case of the V2K), pixels with higher than mean dark current are excluded from the spatial noise calculation. At right, a 20K drop in temperature makes the VIRGO comparable to the H2RG, provided that only pixels in the lower half of the dark current histogram are included in the noise calculation.

A single 1Kx1K InGaAs array was manufactured by Goodrich (formerly Sensors Unlimited) and hybridized at Rockwell to a H1RG multiplexor. Performance, described by Seshadri¹⁰, was commendable for a first attempt but this technology was not pursued due to the funds deemed necessary to bring large format devices to the same level of maturity as the HgCdTe candidates. The noise performance was comparable to the other detectors at low temperatures but was degraded by dark current in 300s exposures at 140K.

14. CONCLUSION AND PLANS

A range of noise performances have been seen in Rockwell arrays, with two (H2RG-102 and 103) meeting the noise and dark current requirements, while simultaneously delivering very high QE, smooth intrapixel response. The substrate has been removed from these devices so that particle events have minimal impact. While the spectrograph could benefit from further noise reduction in long exposures with deep multiple sampling, the current noise is not unacceptable.

Further development at Rockwell and Raytheon has been funded by the SNAP project, which will soon produce new devices from both vendors. Rockwell has advanced their understanding of process parameters, which will lead to further noise improvements and/or increased yield. Dark current improvements hoped for in Raytheon devices will allow the more detailed characterization of their noise. A dual vendor procurement for SNAP's very large NIR focal plane remains likely.

Ongoing lab work will be directed towards further understanding of noise, persistence and dependence of sensitivity and linearity on flux and exposure duration, which – if early indications are correct – may be different manifestations of similar charge capture/emission mechanisms on a variety of timescales.

ACKNOWLEDGEMENT

This work was supported by Department Of Energy grant DE-FG02-04ER41316.

REFERENCES

1. M. Schubnell et al., *Near infrared detectors for SNAP*, Astronomical Telescopes and Instrumentation 2006, Proc. SPIE 6276 paper 29.
2. Matthew G. Brown, *Development of NIR Detectors and Science Driven Requirements for SNAP*, Astronomical Telescopes and Instrumentation 2006, Proc. SPIE 6265 paper113.
3. J. Janesick, J. Andrews, T. Elliott, *Fundamental performance differences between CMOS and CCD imagers: Part I*, Astronomical Telescopes and Instrumentation 2006, Proc. SPIE 6276 paper 77.
4. A. Moore, Z. Ninkov, W. Forrest, *Interpixel capacitance in non-destructive focal plane arrays*, Proc. SPIE Volume 5167 (2003)
5. A. Moore, Z. Ninkov, W. Forrest, *QE Overestimation and Deterministic Crosstalk Resulting from Interpixel Capacitance*, Optical Engineering, Volume 45-7
6. G. Finger, J.W. Beletic, R. Dorn, M. Meyer, L. Mehrgan, A. Moorwood, J. Stegmeier, *Conversion Gain and Interpixel Capacitance of CMOS Hybrid Focal Plane Arrays*, Proc. Scientific Detectors for Astronomy 2005, J.E. Beletic, J.W. Beletic, P. Amico eds., Springer Astrophysics and Space Sciences Library Vol .336, 477-490.
7. D. Hall, *MBE MCT Arrays For JWST: Ultra Low Background Characterization and Temperature Drift Compensations*, Scientific Detectors for Astronomy 2005, J.E. Beletic, J.W. Beletic, P. Amico eds., Springer Astrophysics and Space Sciences Library Vol .336, 491-498.
8. R. Smith, M. Bonati, D. Guzman, *VIRGO-2K 2.5 μ m HgCdTe Dark Current*, Astronomical Telescopes and Instrumentation 2004, Proc. SPIE Volume 5499, 119-130.
9. M. Loose, Rockwell Scientific, *Sensitivity Analysis for the HAWAII-2RG multiplexer*, private communication (2004)
10. S. Seshadri, D. Cole, R. Smith *Characterization of NIR InGaAs imager arrays for the JDEM SNAP mission concept*, Astronomical Telescopes and Instrumentation 2006, Proc. SPIE 6276 paper 31.